arXiv:1704.03958v1 [cs.CV] 13 Apr 2017

Efficient Sparse Subspace Clustering by Nearest Neighbour Filtering

Stephen Tierney[a,*], Yi Guo[b], Junbin Gao[a]

[a] *Discipline of Business Analytics, The University of Sydney Business School, The University of Sydney, NSW 2006, Australia*
[b] *Centre for Research in Mathematics, School of Computing, Engineering and Mathematics, Western Sydney University, Parramatta, NSW 2150, Australia*

Abstract

Sparse Subspace Clustering (SSC) has been used extensively for subspace identification tasks due to its theoretical guarantees and relative ease of implementation. However SSC has quadratic computation and memory requirements with respect to the number of input data points. This burden has prohibited SSCs use for all but the smallest datasets. To overcome this we propose a new method, k-SSC, that screens out a large number of data points to both reduce SSC to linear memory and computational requirements. We provide theoretical analysis for the bounds of success for k-SSC. Our experiments show that k-SSC exceeds theoretical expectations and outperforms existing SSC approximations by maintaining the classification performance of SSC. Furthermore in the spirit of reproducible research we have publicly released the source code for k-SSC[1]

1. Introduction

As the resolution of capture devices continue to increase so does the burden on analytical and classification algorithms. Furthermore high dimensional data is subject to the curse of dimensionality. It is thus necessary to reduce the dimensionality of data to facilitate data analysis. Most common dimension

*Corresponding author
Email addresses: stephen.tierney@sydney.edu.au (Stephen Tierney), y.guo@westernsydney.edu.au (Yi Guo), junbin.gao@sydney.edu.au (Junbin Gao)

[1] https://github.com/sjtrny/kssc

reduction is performed using Principal Component Analysis (PCA) [1, 2, 3], which takes a collection of data points from their original high dimensional space and fits them to a lower dimensional subspace. PCA and associated techniques assume that the entire dataset occupies a single subspace. However in reality large datasets are often composed of a union of subspaces. Therefore it is imperative that the individual subspaces are identified so that the data can be partitioned and dimension reduction can be performed on each subspace separately. The task of assigning each data point to its respective subspace is known as subspace clustering.

More formally we express the subspace clustering problem as follows: given a data matrix of N observed column-wise samples $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N] \in \mathbb{R}^{D \times N}$, where D is the dimension of the data, the objective of subspace clustering is to learn the corresponding subspace labels $\mathbf{l} = [l_1, l_2, \ldots, l_N] \in \mathbb{N}^N$ for all the data points where each $l_i \in \{1, \ldots, p\}$. Data within $\mathbf{X}$ is assumed to be drawn from a union of p subspaces $\{S_i\}_{i=1}^p$ of dimensions $\{d_i\}_{i=1}^p$. Both the number of subspaces p and the dimension of each subspace d_i are unknown. To further complicate the problem it is rarely the case that $\mathbf{X}$ is noise or corruption free. The data is often subject to noise or corruption either at the time of capture (e.g. a digital imaging device) or during transmission (e.g. wireless communication). It is quite clear that subspace clustering is a difficult task since one must produce accurate results quickly while contending with numerous unknown parameters and large volume of potentially noisy data.

[4] introduced an elegant method for subspace clustering called "Sparse Subspace Clustering" (SSC). SSC exploits the self-expressive property of data [5] to find the subspaces:

> *each data point in a union of subspaces can be efficiently reconstructed by a combination of other points in the data*

which gives the relation

$$\mathbf{x}_i = \mathbf{X}\mathbf{z}_i, \tag{1}$$

where $\mathbf{z}_i$ is a vector of reconstruction coefficients for $\mathbf{x}_i$. We can then construct a model for the entire dataset as $\mathbf{X} = \mathbf{XZ}$. In this unrestricted case there are near infinite possibilities for the coefficient matrix $\mathbf{Z}$. Fortunately to reconstruct each data point one only needs $\{d_i\}$ other points. This means that each data point can be sparsely represented by the other points. Sparse Subspace Clustering as its name suggests exploits this fact. The objective function for SSC is

$$\min_{\mathbf{Z}} \|\mathbf{Z}\|_1 \quad \text{s.t. } \mathbf{X} = \mathbf{XZ}, \text{diag}(\mathbf{Z}) = \mathbf{0}, \tag{2}$$

where $\|\mathbf{Z}\|_1 = \sum_i \sum_j |Z_{ij}|$ (the ℓ_1 norm) is used as a surrogate for the ℓ_0 norm and the diagonal constraint prevents the data point from being represented by itself.

Solving such an objective is only useful when the data $\mathbf{X}$ is known to be noise free. As previously mentioned this is extremely unlikely in practice. SSC and most other subspace clustering methods assume the following data generation model

$$\mathbf{x}_i = \mathbf{a}_i + \mathbf{n}_i. \tag{3}$$

where $\mathbf{a}_i$ is the latent original data vector and $\mathbf{n}_i$ is Gaussian noise. Fortunately it is not necessary to recover the original data to perform subspace clustering. To overcome this one can extend the objective to take noise into consideration by relaxing the constraint

$$\min_{\mathbf{Z}} \lambda\|\mathbf{Z}\|_1 + \frac{1}{2}\|\mathbf{X} - \mathbf{XZ}\|_F^2 \quad \text{s.t. } \text{diag}(\mathbf{Z}) = \mathbf{0}. \tag{4}$$

Alternatively one can instead solve an exact variant of the objective by incorporating a fitting error term $\mathbf{E}$ i.e.

$$\min_{\mathbf{Z}} \lambda\|\mathbf{Z}\|_1 + \frac{1}{2}\|\mathbf{E}\|_F^2 \quad \text{s.t. } \mathbf{X} = \mathbf{XZ} + \mathbf{E}, \text{diag}(\mathbf{Z}) = \mathbf{0} \tag{5}$$

where λ is used to control the sparsity of $\mathbf{Z}$. Implementation details for relaxed and exact SSC can be found in Appendices Appendix C and Appendix D respectively.

To obtain the final subspace labels the reconstruction coefficients in $\mathbf{Z}$ are given a secondary interpretation as the affinity or similarity between the data points. Spectral clustering is applied to $\mathbf{Z}$. Typically N-CUT [6] is used as it produces the most accurate segmentation even for poorly constructed affinity matrices and is relatively fast.

While SSC has promising theoretical guarantees [5] and has shown great performance for small evaluation datasets it is not widely applied in practice. This is due to the following:

- $\mathrm{O}\big(N^2\big)$ memory requirements;
- $\mathrm{O}\big(N^2\big)$ flop (floating point operations) requirements.

The first is easily understood as $\mathbf{Z} \in \mathbb{R}^{N \times N}$. One could contend that $\mathbf{Z}$ could be stored in a sparse format, however since the support of $\mathbf{Z}$ is unknown and varies between iterations of the SSC algorithm this approach would introduce significant overhead. Similarly the high flop count is due to the dimensions of $\mathbf{Z}$, since each element must be calculated per iteration.

Our Contributions: In this paper we propose a new algorithm called k-SSC, which is designed for big-data applications. k-SSC dramatically reduces the memory requirements and computation time compared with pre-existing algorithms. Furthermore the conditions of correct subspace identification are provided through theoretical analysis. The rest of the paper is structured as follows: In Section 2, we further discuss the subspace clustering problem and provide an overview of related work. Section 3 is dedicated to discussing the motivation for k-SSC along with its operation and a brief sketch of the theoretical clustering analysis. Detailed proofs for the theoretical analysis are provided in Appendix Appendix A. We follow this with an optimisation scheme and complexity analysis. Next in Section 4 we provide empirical analysis of kSSC using synthetic data. We then finish with a collection of experiments on real world data in Section 5 and some concluding remarks in Section 6.

2. Background and Related Work

The union of subspaces model is applicable to a wide variety of data and thus exploited for a large number of applications. Examples include identifying individual rigidly moving objects in video [7, 8, 9, 10], identifying face images of a subject under varying illumination [11, 12], image compression [13], image classification [14, 15], feature extraction [16, 17], image segmentation [18, 19], segmentation of human activities [20], temporal video segmentation [21, 22] and segmentation of hyperspectral mineral data [23, 21]

This huge range of applications for subspace clustering has spurred the development of subspace clustering algorithms. Early algebraic methods such as Generalised Principal Component Analysis (GPCA) [22, 24] suffered from sensitivity to noise and increasing computational complexity as the number and size of subspaces increases [5]. A number of statistical methods have been developed such as Mixtures of Probabilistic PCA (MPPCA) [25] and Multi-Stage Learning (MSL) [26] that make Gaussian assumptions about the distribution of the data in the subspace. However such approaches rely on good initialisation and the dimensions of the subspaces must be known before hand.

More recently spectral methods have come to dominate subspace clustering literature as they offer a vast improvement over the previously mentioned methods. First they do not increase in complexity with the number or dimension of subspaces, second they are robust to noise or outliers and finally they provide a simple to understand work pipeline that is easily adapted and modified. Spectral methods consist of two stages: learning a similarity matrix for the data then assignment of class labels through segmentation of the similarity matrix. We have already introduced SSC, the forerunner of spectral subspace clustering. Other spectral methods adopt the general SSC objective but impose different structural constraints over the similarity matrix $\mathbf{Z}$ instead of the ℓ_1 norm. For example Low-Rank Representation (LRR) [27] imposes a rank penalty to obtain a more globally consistent $\mathbf{Z}$. Or different penalties are used to minimise fitting error depending upon the type of expected noise and whether or not outliers

are likely [28, 29]. Further regularisation can be applied to incorporate prior knowledge such as the spatial structure or sequence of the data [23, 21, 30].

Another factor that has spurred on research of spectral methods is that they can be guaranteed to correctly segment the subspaces. For example it was shown by [31] that correct subspace identification is guaranteed for SSC provided that data driven regularisation for each λ_i (column wise splitting of $\mathbf{Z}$) is used. This of course is subject to further conditions such as a minimum distance between subspaces, sufficient sampling of points from each subspace and the noise level in the data. For LRR the requirements for guaranteed success are much stricter, please see [27] for full details.

However despite the aforementioned strengths of spectral clustering algorithms they suffer from huge memory or computational requirements that prevent them from being applied to even modestly sized datasets. In the academic literature this problem is generally ignored as the data used for analysis is very limited in the number of data points. In light of these issues there has been considerable interest in developing tractable subspace clustering algorithms. Unfortunately they either lack theoretical justification or suffer dramatically in terms of clustering accuracy in practice.

SSC approximation methods can be divided into two classes: inductive and heuristic. Inductive methods perform SSC or learn the similarity matrix on a small subset of the data. This full structure of the similarity matrix or labels is then obtained by inductive transfer from the subset. Heuristic methods abandon SSC entirely and directly assign class labels by greedy selection of nearest neighbours based on a defined metric.

Scalable SSC (SSSC) [32] was the first attempt to resolve computational issues. As an inductive method it first selects some candidate samples from the data and performs SSC on these samples. Then the remaining samples are assigned to clusters based on their fit into the clusters formed by the training candidates. This approach has considerable issues. First the candidate samples are selected by uniform random sampling. This does not guarantee that every cluster will be accounted for in the candidate set. Second there must be enough

candidate samples for the candidate clusters to generalise to the remaining samples. Correctly choosing the number of samples is a difficult task.

Arguably the most prominent Heuristic method is Orthogonal Matching Pursuit (OMP), which has been long used as a greedy sparse approximation method [33]. For each data point a residual vector is set as the data point. Then the nearest neighbour to the residual is found and then the residual is updated by a projection of the data point onto the span formed by the currently picked up neighbours. This is repeated until the number of neighbours is reached or the norm of the residual is small enough. OMP is also known by other names such as Greedy Feature Selection (GFS) [34] and is a constant well that researches draw from [35].

Although OMP is advertised as being a fast approximate method however in practice we do not find this to be the case. First the nearest neighbour search is performed for every iteration. Second the computational and memory requirements of a Naive implementation tend to increase dramatically as D increases due to the need to create a $D \times D$ matrix in each iteration for every data point. This, in some cases, makes it just as intractable as SSC. Third the Naive implementation requires successive computation of the SVD of the span matrix at each iteration. The second and third points have fortunately been mitigated by improvements such as Rank-1 updating scheme of Moore-Penrose Inverse [36], factorisation approaches such as QR and Cholesky decomposition [37] or more esoteric methods [33, 38, 39].

However these accelerations to the OMP algorithm are relatively meaninginglyess as we find the reported accuracy results of methods using OMP to be dubious. As shown in Section 5 we find that GFS (OMP) performs poorly in terms of clustering accuracy. We stress that we were unable to reproduce the promising results of OMP based methods as claimed by earlier works.

OMP has inspired other methods such as Greedy Subspace Clustering (GSC) [40] and ORGEN [41]. GSC differs from OMP in neighbour selection. Each neighbour is selected by finding the data point which has the largest norm of projection onto the span formed by the current neighbours. Although at

first glance GSC appears to be simple and thus likely to scale well w.r.t. N, the projection step is quite computationally intensive and just like OMP the nearest neighbour search is performed in each iteration.

ORGEN extends the SSC model to the Elastic-Net model. That is, it uses the ℓ_2 norm in tandem with the ℓ_1 norm. From an initialisation point of some neighbours of each $\mathbf{x}_i$ it solves the Elastic-Net objective then determines an "oracle point", which is the residual from fitting the coefficients from the Elastic-Net procedure to the model. This oracle point is then used to find potential new neighbours. The procedure terminates when no new neighbours are added at the end of an iteration. ORGEN suffers from a number of problems. First it is highly initialisation dependant. The authors suggest performing ℓ_2 sub problem and choosing the largest valued elements as the initialisation pool for each $\mathbf{x}_i$. This can be slow as when D is large and the ℓ_2 problem lacks rigorous guarantees of successful subspace identification. Second the repeated computation of elastic net is problematic when the active set grows large. This is a very real concern as termination only occurs when the active set stops growing. The active set could grow to the full size of the data set. Third the claim of improved running time is not evident. The authors show running times for single $\mathbf{x}_i$ instead of the whole data $\mathbf{X}$ and do not compare to different approaches such as [32] or [42].

Heuristic methods are often incredibly simple. For example Robust Subspace Clustering via Thresholding (TSC) [42] essentially performs nearest neighbour based spectral clustering. For each point the nearest neighbours are found and the affinity matrix is constructed using exponential inner product between each of the neighbours.

3. Efficient Sparse Subspace Clustering

Our contribution to subspace clustering is inspired by the sparsity of SSC. It has been shown repeatedly that each data point can be reconstructed by only d_i (the dimensionality of the underlying subspace) other data points from its corresponding subspace [5, 31, 43]. This is the basis of SSC's operation. By

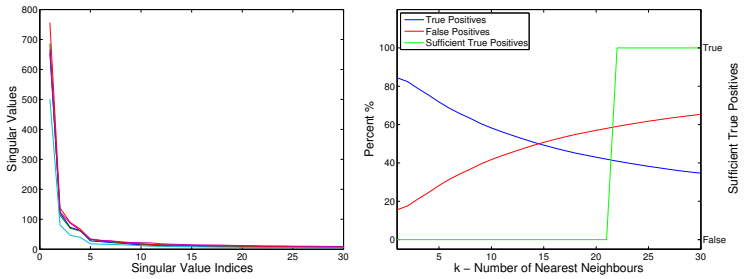

Figure 1: Left: Singular values of several faces from the Extended Yale B dataset. Right: Average percentage of true positive and false positive nearest neighbour selection from the entire Extended Yale B dataset and correspondingly, in green, a plot of when the sufficient true positives are reached on average.

finding the sparsest representation one will be left with the minimum support to represent a data point $\mathbf{x}_i$, corresponding to data points in the same subspace. Therefore it is clear that blindly considering every data point as a candidate for reconstruction is very wasteful since only a relative few points will be left as support. Furthermore the process is very intensive in terms of memory requirements. The algorithm for solving the exact variant of SSC (see Appendix Appendix D) requires $\mathrm{O}\big(N^2\big)$ FLOPs per iteration and the storage of $\mathrm{O}\big(N^2\big)$ floats over the algorithm's entire operation w.r.t. $\mathbf{Z}$.

Therefore if we can safely prune a vast majority of data points as candidates for another data points reconstruction then we can massively reduce computational and memory load. In other words this means that we would only solve for a small subset of the entries of $\mathbf{Z}$ rather than the entire matrix. To this end we propose kSSC, in which we limit each data point to be represented by at most k other data points. Thus the relaxed objective function for kSSC is

$$\min_{\mathbf{Z}} \lambda\|\mathbf{Z}\|_1 + \frac{1}{2}\sum_{i}^{N} \|\mathbf{x}_i - \sum_{j\in\Omega_i} \mathbf{x}_j Z_{ji}\|_F^2 \quad \text{s.t. } \mathrm{diag}(\mathbf{Z}) = \mathbf{0}. \tag{6}$$

where Ω_i is the set of data points to use for reconstruction of data point i. Under this objective we can reduce both the memory and FLOP requirements to $\mathrm{O}\big(kN\big)$ w.r.t. $\mathbf{Z}$, which when $k \ll N$ provides massive savings.

Evidently the success of kSSC relies heavily upon both the size of k and the

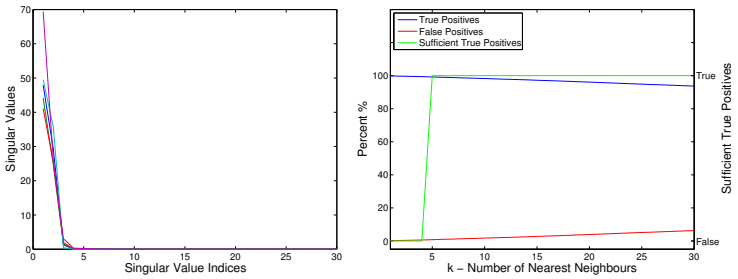

Figure 2: Left: Singular values of several motions from the Hopkins 155 Motion Dataset. Right: Average percentage of true positive and false positive nearest neighbour selection from the checkerboard and traffic sequences in the Hopkins 155 Motion Dataset dataset and correspondingly, in green, a plot of when the sufficient true positives are reached on average.

scheme that is used to select Ω_i. First one should always choose $k \geq d_i$ since each data point needs at a minimum d_i other points for reconstruction. The dimension of a subspace can be roughly estimated by analysing the singular values of the matrix of samples from a subspace. For example in Figure 1 we show singular values for multiple subspaces (a single subspace corresponds to a single subject or face) from the Extended Yale B dataset. The point at which the singular values begin to trail off reveals the underlying subspace dimension d_i, which in this case is 9 [5]. Therefore in that case we must set $k \geq 9$. Similarly in Figure 2 we perform the same analysis on the motion segmentation dataset and find that the subspace dimension is 4. Since in almost all cases $d_i \ll N$ and thus $k \ll N$ the computational and memory requirements of kSSC will be much lower than SSC. However even in cases where there is no ground truth or sample data available one can set a large, conservative value for k with little impact on overall performance. For example in Figure 3 we show that increasing k from 100 to 1000 barely impacts FLOP count or memory requirements relative to the original requirements of SSC. These values were calculated using Table 1.

Second one must choose Ω_i such that it contains enough data points from $\mathbf{x}_i$'s subspace. Uniformly random sampling to choose k neighbours is a poor choice since the selected neighbours may not belong to the same subspace. Recent works such as [34, 40, 43] have demonstrated that even in noisey cases or cases

of subspace intersection that the points closest to each $\mathbf{x}_i$ in the ambient space usually correspond to the most strongly connected data points in $\mathbf{Z}$ i.e. data points from the same subspace. We come to the same conclusions in Appendix Appendix A where we prove that in both noisey and noise free cases we are able to correctly select points from the subspace using k nearest neighbours. We repeat our two central theorems here for the reader. First in the case of noise free data:

Theorem 3.1 (Recalled from Theorem Appendix A.5)**.** *Let d_m be the minimum dimensionality of all the subspaces. Given the conditions in Theorem Appendix A.2, if $\mathcal{A}_{\ell,1} \leq \frac{\sqrt{d_m}(1-\epsilon^2/2)}{2\sqrt{d(t\log N_\ell + t^2)}}$, then the samples selected for any sample $\mathbf{x}_1$ from subspace $\mathcal{S}_1$ by using kNN (with $k = k_0/C$) contains no samples from other subspaces but $\mathcal{S}_1$ with probability at least $1 - 2e^{-t} - e^{-k_0}(eC)^{k_0/C}$.*

Second in the case of noisey data

Theorem 3.2 (Recalled from Theorem Appendix A.9)**.** *Let d_m be the minimum dimensionality of all the subspaces. Given the conditions in Theorem Appendix A.2 and the noise model* (A.3)*, for a small positive δ, if $\mathcal{A}_{\ell,1} \leq \frac{\sqrt{d_m}(1-\epsilon^2/2-6\delta)}{2\sqrt{d(t\log N_\ell + t^2)}}$, then the samples selected for any sample $\mathbf{y}_1$ from subspace $\mathcal{S}_1$ by using kNN (with $k = k_0/C$) contains no samples from other subspaces but $\mathcal{S}_1$ with probability at least $1 - 2e^{-t} - e^{-k_0}(eC)^{k_0/C} - 2\exp(1 - \frac{c\delta^2}{\sigma^2}) - \frac{d\sigma^4}{\delta^2}$.*

Therefore we set Ω_i for $\mathbf{x}_i$ as its k nearest neighbours from the original ambient space. However when $\mathbf{X}$ is subject to noise this assumption no longer holds. For this reason we recommend setting k well above d_i to provide sufficient head room. Furthermore we suggest increasing k as the magnitude of expected noise increases, since as noise increases, so does the likelihood of false positive neighbour selection. In Figure 1 and Figure 2 we demonstrate this effect on the Yale and Rigid Motion datasets respectively. We note that the required value for k to select sufficient true positives via kNN exceeds d_i. This is due to the presence of noise and corruptions in the data and the sometimes small distance between subspaces, particularly for the Extended Yale B dataset. Although still extremely small relative to N.

Algorithm 1 kSSC

Require: $\mathbf{X}^{D\times N}$ - observed data, k - number of neighbours, c - number of subspaces

1: **for** $i \to N$ in parallel **do**
2: Set Ω_i by kNN
3: Obtain coefficients $\mathbf{z}_i$ by solving (6)
4: **end for**
5: Form the similarity graph $\mathbf{W} = |\mathbf{Z}| + |\mathbf{Z}|^T$
6: Apply N-Cut to $\mathbf{W}$ to partition the data into c subspaces
7: **return** Subspaces $\{S_i\}_{i=1}^{c}$

In summary we propose to eliminate the calculation of redundant elements of $\mathbf{Z}$ by computing only k rather than N coefficients for each $\mathbf{x}_i$. An overview of the entire method can be found in Algorithm 1. Subspace identification accuracy can be exactly maintained from SSC provided that the following conditions are met:

- k is equal to or greater than $\max(d_i)$
- the elements of Ω_i are nearest neighbours of $\mathbf{x}_i$

These conditions are sufficient but not necessary. In some cases clustering accuracy could be maintained when k is less than $\max(d_i)$ or different filtering method is used. However when these conditions are met kSSC ensures that SSC's guarantee of correct subspace identification and robustness to noise is preserved since kNN is guaranteed to correctly identify neighbours (see Appendix Appendix A). Furthermore kSSC is easily solved in parallel as Ω_i and each column of $\mathbf{Z}$ is independent.

3.1. Optimisation

To solve (6) we use FISTA (Fast Iterative Shrinkage-Thresholding Algorithm) [44, 45]. FISTA is an accelerated gradient descent scheme for solving objective functions containing a smooth part and non-smooth part as is the case

with (6). One of the key abilities of FISTA is that it guarantees a convergence rate of $\mathrm{O}\big(\frac{1}{\mathtt{t}^2}\big)$ where $\mathtt{t}$ is the iteration counter. This is achieved by dynamically setting the rate of descent parameter (Lipschitz constant) and using the two previous iteration points to accelerate the gradient descent. Furthermore FISTA provides the aforementioned ability with minimal computational and memory overhead. Each iteration of FISTA only requires solving a closed form proximity problem which in the case of ℓ_1 minimisation can be solved at an element wise level. This allows us to resolve the selective fitting term of (6) since we can enforce it by ignoring the elements of $\mathbf{Z}$ that are outside of Ω.

We begin by re-writing, with some abuse of notation, the original objective (6) for a single column of $\mathbf{Z}$

$$\min_{\mathbf{z}_i} L = \lambda\|\mathbf{z}_i\|_1 + \frac{1}{2}\|\mathbf{x}_i - \mathbf{X}_i\mathbf{z}_i\|_2^2 \tag{7}$$

where $\mathbf{z}_i = \mathbf{z}_{\Omega_i i}$ i.e. the vector of rows Ω_i and column i of $\mathbf{Z}$ and $\mathbf{X}_i = \mathbf{X}_{(:,\Omega_i)}$ i.e. the matrix formed from the columns of $\mathbf{X}$ indexed by Ω_i. Note that we have removed the constraint $\mathrm{diag}(\mathbf{Z}) = \mathbf{0}$ since we enforce it by ensuring that no diagonal entries are present in each Ω_i.

At each iteration in the FISTA scheme one must solve the ℓ_1 proximal linearised form of L. Denote the linearisation of L at point $\mathbf{z}_i^{\mathtt{t}}$

$$\min_{\mathbf{z}_i} \widetilde{L}_\rho(\mathbf{z_i}, \mathbf{z_i^t}) = \lambda\|\mathbf{z}_i\|_1 + \frac{\rho}{2}\|\mathbf{z}_i - (\mathbf{z}_i^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{z}_i^{\mathtt{t}}))\|_2^2, \tag{8}$$

where $F = \frac{1}{2}\|\mathbf{x}_i - \mathbf{X}_i\mathbf{z}_i\|_2^2$ and correspondingly $\partial F = -\mathbf{X}_i^T(\mathbf{x}_i - \mathbf{X}_i\mathbf{z}_i^{\mathtt{t}})$. The solution to (8) is given by the closed-form ℓ_1 shrinkage function $\mathcal{S}_\tau$ as follows

$$\mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{z}_i^{\mathtt{t}}) = \mathrm{sign}(\mathbf{z}_i^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{z}_i^{\mathtt{t}}))\max(|\mathbf{z}_i^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{z}_i^{\mathtt{t}})| - \frac{\lambda}{\rho}, 0). \tag{9}$$

We refer readers to [46, 47] for further details. The full algorithm is outlined in Algorithm 2.

3.2. Segmentation

After solving (6) for each $\mathbf{z}_i$ the next step is to form $\mathbf{Z}$ and use the information encoded in $\mathbf{Z}$ to assign each data point to a subspace. A robust approach is

Algorithm 2 Solving (6) via FISTA

Require: $r_i = \infty$, $\mathbf{z}_i = \mathbf{0}$, $\mathbf{j}_i = \mathbf{0}$, $\alpha_i = 1$, λ, ρ_i, γ, ϵ

  while $r_i^{\mathtt{t}} - r_i^{\mathtt{t}-1} \geq \epsilon$ **do**

    while $L(\mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{j}_i^{\mathtt{t}})) \geq \widetilde{L}_\rho(\mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{j}_i^{\mathtt{t}}), \mathbf{j}_i^{\mathtt{t}})$ **do**

      $\rho_i = \gamma \rho_i$

    end while

    $\mathbf{z}_i^{\mathtt{t}+1} = \mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{j}_i^{\mathtt{t}}))$

    $\alpha_i^{\mathtt{t}+1} = \frac{1+\sqrt{1+4\alpha_i^{t2})}}{2}$

    $\mathbf{j}_i^{\mathtt{t}+1} = \mathbf{z}_i^{\mathtt{t}+1} + \left(\frac{\alpha_i^{\mathtt{t}}-1}{\alpha_i^{\mathtt{t}+1}}\right)(\mathbf{z}_i^{\mathtt{t}+1} - \mathbf{z}_i^{\mathtt{t}})$

    $r_i^{\mathtt{t}+1} = L(\mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{j}_i^{\mathtt{t}}))$

  end while

to use spectral clustering. The matrix $\mathbf{Z}$ can be interpreted as the affinity or distance matrix of an undirected graph. Element Z_{ij} corresponds to the edge weight or affinity between vertices (data points) i and j. Then we use the spectral clustering technique, Normalised Cuts (N-Cut) [6], to obtain final segmentation. N-Cut has been shown to be robust in subspace segmentation tasks and is considered state of the art. Since we expect $\mathbf{Z}$ to be sparse in most cases N-Cut should have reasonable computation time, particularly in comparison to a full $\mathbf{Z}$ matrix. However in cases where N-Cut is too slow one can use approximate techniques such as the Nyström method [48]. Spectral segmentation techniques such as N-Cut require the number of subspaces p as a parameter. In this paper we assume that the number of subspaces is estimated by the user although automatic techniques exist for the estimation of p, see [21] for details.

Spectral segmentation techniques such as N-Cut require the number of subspaces p as a parameter. In the case where the number of subspaces is unknown one can use either the Eigen-gap [49, 50, 31] or the closely related SVD-gap heuristic of [27]. The Eigen-gap heuristic uses the eigenvalues of $\mathbf{W}$ to find the number subspaces. It does this by finding the largest gap between the ordered eigenvalues, the number of eigenvalues before this point is treated as the number

Method	FLOPs (CPU)	Floats (RAM)
SSC Exact	$7N^2 + 4DN^2 + 4DN$	$2N^2 + DN$
SSC Relaxed	$6N^2 + 4DN^2 + DN$	$4N^2$
kSSC Exact	$7kN + 4kDN + 4DN$	$2kN + DN$
kSSC Relaxed	$6kN + 4kDN + DN$	$4kN$

Table 1: Overview of FLOP requirements w.r.t $\mathbf{Z}$ per iteration of SSC and kSSC and Float requirements over entire operation of SSC and kSSC.

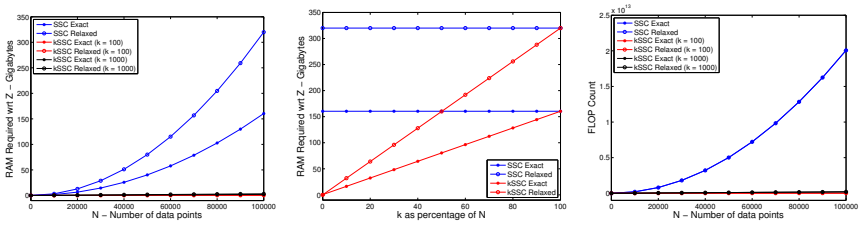
SSC Exact
SSC Relaxed
kSSC Exact (k = 100)
kSSC Relaxed (k = 100)
kSSC Exact (k = 1000)
kSSC Relaxed (k = 1000)
RAM Required wrt Z – Gigabytes
N – Number of data points
k as percentage of N
FLOP Count

(a) Memory requirements w.r.t. $\mathbf{Z}$ for SSC and kSSC. $D = 500$.

(b) Memory requirements w.r.t. $\mathbf{Z}$ for kSSC as $k \to N$. $D = 500$.

(c) FLOP requirements w.r.t. $\mathbf{Z}$ for a single iteration of SSC and kSSC. $D = 500$.

Figure 3: Visual comparison of memory and computation requirements from Table 1. Even as k grows much larger than d_i it has little impact on computation time and memory.

of clusters. Let $\{\delta_i\}_{i=1}^N$ be the descending sorted eigenvalues of $\mathbf{W}$ such that $\delta_1 \geq \delta_2 \geq \cdots \geq \delta_N$. Then p can be estimated by

$$p = \underset{i=1,\ldots,N-1}{\operatorname{argmax}} \ (\delta_i - \delta_{i+1})$$

The SVD-gap heuristic is the same procedure with eigenvalues of $\mathbf{W}$ replaced with singular values. Further improvements upon the Eigen-gap heuristic have been made, see [49] for details.

3.3. Complexity Analysis

The complexity of kSSC only varies from SSC w.r.t. $\mathbf{Z}$ as can be seen from Algorithm 1 and a comparison of Algorithm 2 and 4. It differs in two ways. First

we must find the k nearest neighbours of each $\mathbf{x}_i$. Fortunately fast approximate methods exist for computing kNN and are freely available in packages such as FLANN [51]. The computation time for kNN is on the order of $\mathrm{O}\big(N \log N\big)$ and $\mathrm{O}\big(\log N\big)$ for preprocessing and searching respectively [52, 53, 54, 51].

Second is the updating of $\mathbf{z}_i$ at each iteration. Since we are only updating k entries of each column of $\mathbf{Z}$ instead of the full N entries the FLOP count is drastically reduced. We enumerate the different FLOP counts per iteration for all columns of $\mathbf{Z}$ in Table 1 and visualise the dramatic improvement that kSSC offers in Figure 3. Similarly the amount of RAM required in the form of floats for updating $\mathbf{Z}$ is drastically reduced. For both FLOP and RAM the complexity is reduced from $\mathrm{O}\big(N^2\big)$ to $\mathrm{O}\big(N\big)$. Note that the relaxed variant has markedly lower FLOP counts than the exact variants. This assumes one execution of the ℓ_1 shrinkage operator per iteration. However in the case of FISTA, a single iteration may require many executions of the shrinkage operator due to the search scheme for optimal rate of descent parameter. If the rate of descent parameter ρ is initialised poorly this may lead to the exact variant solved by LADMPSAP (see Appendix Appendix B) to be quicker. However in practice we find that solving the relaxed variant by FISTA is usually faster since choosing ρ is not a difficult task and can be estimated by running the solver on a small sub section of the data. Furthermore the FISTA based solver will converge much faster than the LADMPSAP solver. We provide a brief sample of running time differences in Figure 4 to illustrate the difference between implementation variants. We also demonstrate the effect of varying the number of available cores for the parallel implementations in Figure 5. We find that in the case of SSC as the number of cores increase the computation time also increases, which indicates that the performance of SSC is not as straight forward as outlined in Table 1. In fact the performance is markedly worse than expected due to the overhead of sharing and multiple accessing of the full data matrix $\mathbf{X}$, which further reinforces the point that SSC does not scale well with large datasets. On the other hand kSSC benefits greatly from increasing the core count and eventually plateaus due to it's own overhead.

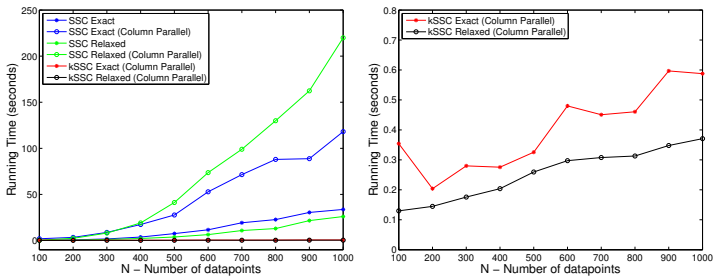

Figure 4: Left: A comparison of running times with increasing N between kSSC, SSC and their various implementations. Right: A zoomed comparison of running times with increasing N for kSSC relaxed and exact variants (taken from the Left plot). Note that the scales are different since kSSC takes a fraction of the running time of SSC.

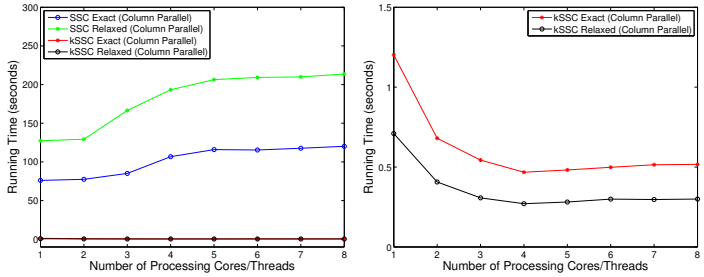

Figure 5: Left: A comparison of running times with increasing threads/cores between kSSC, SSC and their various implementations. Right: A zoomed comparison of running times with increasing threads/core for kSSC relaxed and exact variants (taken from the Left plot). Note that the scales are different since kSSC takes a fraction of the running time of SSC.

4. Preliminary Synthetic Evaluation

In this section we use synthetic data to experimentally evaluate our hypothesis proposed in Section 3 and the therotetical analysis in Appendix Appendix A that kSSC can match the clustering accuracy of SSC.

In an effort to maximise transparency and repeatability, all MATLAB code and data used for these experiments and those in Section 5 can be found online at `https://github.com/sjtrny/kSSC`. To help evaluate consistency parameters except for k were fixed for each experiment, which we further explain in the following subsections and are recorded in the code repository.

4.1. Metrics

Segmentation accuracy was measured using the subspace clustering error (SCE) metric [5], which is defined as

$$\text{SCE} = \frac{\text{num. of misclassified points}}{\text{total num. of points}} \times 100, \tag{10}$$

where lower subspace clustering error means greater clustering accuracy. In cases where we inject extra noise we report the level of noise using Peak Signal-to-Noise Ratio (PSNR) defined as

$$\text{PSNR} = 10 \log_{10} \left(\frac{s^2}{\frac{1}{mn} \sum_i^m \sum_j^n (A_{ij} - X_{ij})^2} \right) \tag{11}$$

where $\mathbf{X} = \mathbf{A} + \mathbf{N}$, $\mathbf{A}$ is the original data, $\mathbf{N}$ is noise and s is the maximum possible value of an element of $\mathbf{A}$. Decreasing values of PSNR indicate increasing amounts of noise.

4.2. Effect of subspace dimension, cluster size and ambient dimension

As noted in other works such as [42, 55] the ratio of the subspace dimension d_i to the number of points in each cluster N_i can play a dramatic role in the clustering accuracy of SSC. However these works also ignore the role of ambient dimension D. In this section we demonstrate the relationship between all three variables.

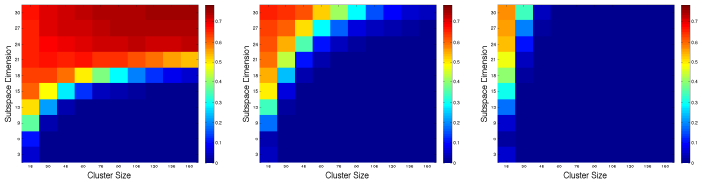

(a) SSC, $D = 30$ (b) SSC, $D = 50$ (c) SSC, $D = 100$

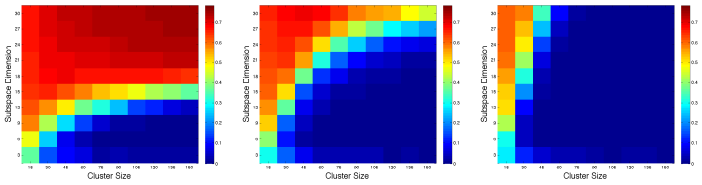

(d) kSSC, $D = 30$ (e) kSSC, $D = 50$ (f) kSSC, $D = 100$

Figure 6: Effect of subspace dimension, cluster size and ambient dimension

We generate 5 subspaces and vary their dimension d_i from 3 to 30 and N_i from 15 to 150. Each subspace is created using random orthonormal vectors as the basis with uniform random coefficients. For each pair of d_i and N_i we take the mean of the SCE over 50 problem instances. We repeat this again for 3 instances with D set to $30, 50$ and 100. For this experiment we set $k = \frac{N_i}{2}$ or $k = 1.5D$, whichever is smaller. The results shown in Figure 6 that kSSC can match the performance of SSC even when $k \ll d_i$.

4.3. Effect of mean, variance and noise in subspace distribution

Previous works have only performed empirical evaluation on uniformly distributed synthetic data. That is, the coefficients chosen are uniform random, as was used in the prior subsection. However we contend instead that in reality most data points encountered in the real world are Gaussian distributed over their respective subspace's basis. Furthermore the data points are often corrupted with noise, which we assume will be $\mathcal{N}(0, 1)$.

For this experiment we vary the mean μ and variance σ^2 of Gaussian distributed data points using random orthonormal vectors as the basis for each

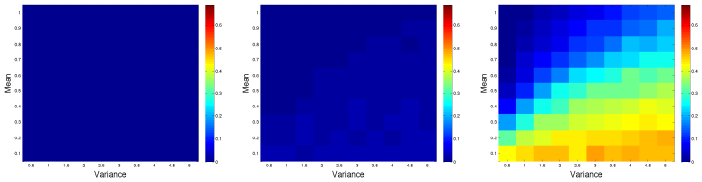

(a) SSC, PSNR 100 (b) SSC, PSNR 60 (c) SSC, PSNR 46

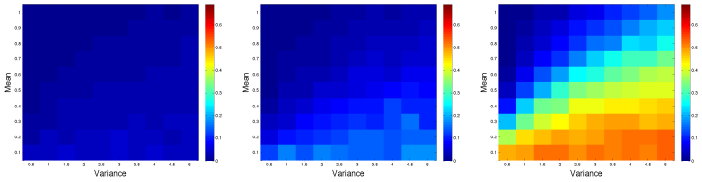

(d) kSSC, PSNR 100 (e) kSSC, PSNR 60 (f) kSSC, PSNR 46

Figure 7: Effect of mean, variance and noise in subspace distribution

subspace. We create 5 subspaces with $d_i = 5$, $N_i = 50$ and $D = 50$. For each pair of μ and σ^2 we take the mean SCE over 50 problem instances. We repeat this again for 3 instances, each time increasing the noise factor, which we report using PSNR. For this experiment we set $k = 10$. The results shown in Figure 7 that kSSC can match the performance of SSC even when $k \ll d_i$. We note that the effect of mean and variance on point distribution in the subspaces is significantly more pronounced as PSNR decreases.

4.4. Effect of subspace intersection

The intersection of subspaces (shared basis vectors) plays an important role in the clustering accuracy. As previously reported by others, the clustering accuracy decreases as the dimension of intersection increases. To demonstrate this effect and that kSSC can match SSC, we perform the same experiment as found in Section 8.1.1 of [42] and Section 5.1.2 of [55]. We generate two subspaces with $D = 200$, $d_i = 10$ and $N_i = 20d_i$ and vary the number of shared basis vectors t from 0 to d_i. We generate $\mathbf{U} \in \mathbb{R}^{D \times 2d_i - t}$ random orthonormal basis vectors and set the basis vectors for $\mathcal{S}_1$ to the first d_i columns of $\mathbf{U}$ and

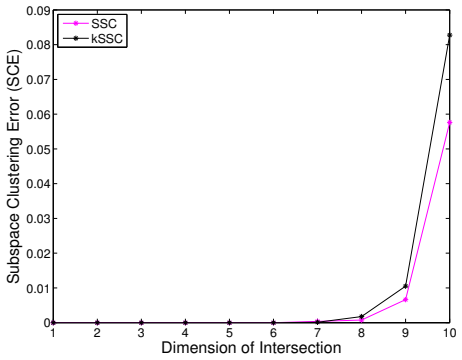

Figure 8: Effect of subspace intersection

correspondingly the basis for $\mathcal{S}_1$ to the last d_i columns. We then take the average SCE over 20 problem instances for each t. Results are reported in Figure 8, where we can clearly see that kSSC closely matches the performance of SSC.

5. Experimental Evaluation

In this section, we evaluate the clustering performance of kSSC on semi-synthetic and real world datasets. We vary the amount of additional noise in some of these experiments to compare the robustness of kSSC against the pre-existing competitor algorithms Greedy Feature Selection (GFS), Greedy Subspace Clustering (GSC), Scalable Sparse Subspace Clustering (SSSC) and Robust Subspace Clustering via Thresholding (TSC). Additionally we use SSC to gauge baseline performance.

The running times of the experiments carried out in Sections 5.1 and 5.3 can be found in Figure 9. Since these experiments are small in size the running time reduction of kSSC is not that significant. However these tests indicate that kSSC matches the clustering accuracy of SSC very closely, an attribute that is not found in other methods. We perform a test in Section 5.2 to evaluate the running time of kSSC for a large scale data set.

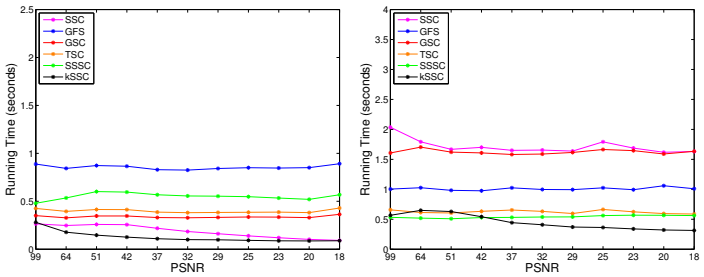

Figure 9: Left: Median running time of each tested algorithm for the Thermal Infrared Data Segmentation experiment found in Section 5.1. Right: Median running time of of each tested algorithm for the Motion Segmentation experiment found in Section 5.3. Overall in these experiments the benefit of kSSC is slight in comparison to other methods since N is low. We refer to readers Section 5.2 for a comparison in running time where N is large.

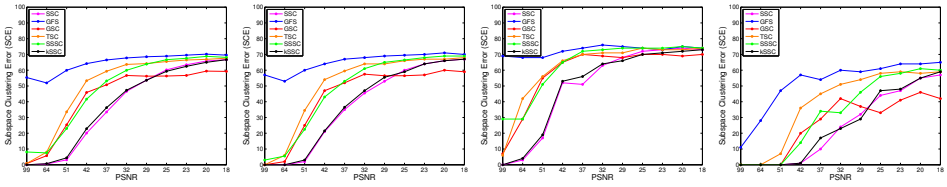

(a) Mean Error (b) Median Error (c) Max Error (d) Min Error

Figure 10: Semi-synthetic Hyperspectral TIR

5.1. Thermal Infrared Data Segmentation

We assemble synthetic data from a library of thermal infrared (TIR) hyper spectral mineral data. The library consists of 120 spectra samples with $D = 321$. We generate 5 subspaces with $d_i = 5$. For each subspace we randomly select 5 basis vectors from the spectra samples in the TIR library and generate 50 points using uniform random coefficients. We then corrupt data with various levels of standard Gaussian noise and evaluate clustering performance of our framework SSC and the Scalable SSC. The experiment is repeated for 50 problem instances for each level of noise to obtain an average SCE. Results can be found in Figure 10. kSSC closely tracks the performance of SSC and outperforms all other methods.

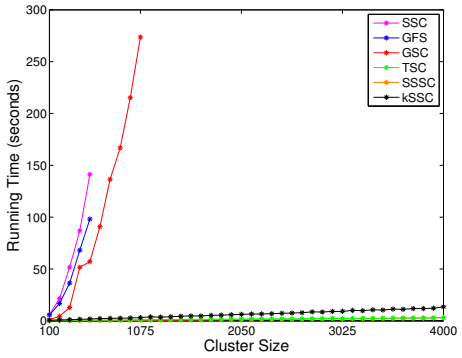

Figure 11: Running time of Large Scale Experiment

5.2. Large Scale Thermal Infrared Segmentation

The main goal of kSSC is to maintain SSC's clustering accuracy but in a fraction of the time. To confirm this ability, we create a large scale semi-synthetic dataset from the TIR data used in the previous subsection. We generate data in a similar fashion to the previous section. However for each subspace we generate N_i points using uniform random coefficients where we vary N_i from 100 to 4000. For this experiment, we stop SSC, GFS and GSC early since they do not scale well in this application (see Section 2). From Figure 11 we find that kSSC has similar run time characteristics to TSC and SSSC.

5.3. Hopkins 155 Motion Segmentation

The aim of this experiment is to assign feature points extracted from a video to their corresponding motion or object in the scene. As previously mentioned, it has been shown in [5] that these features trajectories actually correspond to low-dimensional subspaces. The data from this experiment is drawn from the rigid motion sequences of the Hopkins 155 dataset [56]. These sequences have around 200-500 feature trajectories and range in number of frames from 20-60. Results can be found in Figure 12.. Again kSSC closely tracks the performance of SSC and consistently performs as the PSNR decreases.

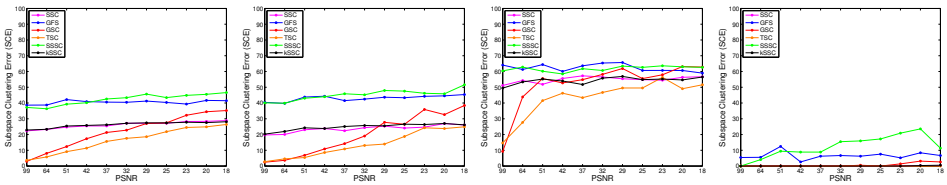

(a) Mean Error (b) Median Error (c) Max Error (d) Min Error

Figure 12: Rigid Motion Segmentation

	Mean	Median	Min	Max	Std	Mean Run Time (s)
SSC	28.1%	31.5%	0.0%	65.6%	24.5%	21.38
GSC	30.6%	29.9%	0.5%	55.2%	15.6%	30.38
TSC	58.1%	61.2%	36.5%	66.1%	7.3%	7.59
SSSC	59.5%	58.6%	31.8%	100.0%	17.2%	0.66
kSSC	22.3%	17.2%	0.0%	65.1%	19.3%	30.38

Table 2: Face Clustering results from the Extended Yale B Dataset.

5.4. Extended Yale B Face Clustering

The aim of this experiment is to cluster unique human subjects from a set of face images. We draw our data from the Exteded Yale Face Database B [12]. The dataset consists of approximately 64 photos of 38 subjects under varying illumination. We select three subjects randomly then resample their images to 96×84 and form data vectors $\mathbf{x}_i \in \mathbb{R}^{2016}$ by concatenating them together. This test was repeated 50 times with new random subjects each time. This is a challenging dataset since the original data is already corrupted by shadows from the varied illumination. Results can be found in Table 2. Surprisingly we find that kSSC even outperforms SSC for this task, which may be due to the aggressive kNN pre-screening process removing all but the most similar data points. In this dataset, there are many face images of different subjects that contain large regions of highly similar data due to the extreme occlusions from shadows. We believe the nearest neighbour filtering selection helps to prevent the possibility of extreme false positives connections in $\mathbf{Z}$.

6. Conclusion

In this paper we proposed a new algorithm, kSSC, to accurately and tractably approximate SSC for large scale datasets. By accurately screening out the vast majority of eligible data points as neighbours the memory and computational requirements are reduced from $\mathrm{O}(N^2)$ to $\mathrm{O}(N)$. Our theoretical analysis shows that we can theoretically match the subspace identification performance of SSC provided that we have sufficient sampling and the magnitude of the noise is not too great. Moreover our empirical results on synthetic and real data demonstrate that kSSC outperforms the existing SSC approximation methods in terms of accuracy and matches or beats the computational and memory requirements.

Acknowledgment

The research project is supported by the Australian Research Council (ARC) through grant DP140102270. It is also partly supported by the Natural Science Foundation of China (NSFC) through project 41371362.

References

References

[1] I. Jolliffe, Principal Component Analysis, Wiley Online Library, 2002.

[2] J. Ramsay, B. W. Silverman, Functional Data Analysis, Springer Series in Statistics, Springer, 2005.

[3] J. Jacques, C. Preda, Functional data clustering: a survey, Advances in Data Analysis and Classification 8 (3) (2014) 231–255.

[4] E. Elhamifar, R. Vidal, Sparse subspace clustering, in: IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 2790–2797.

[5] E. Elhamifar, R. Vidal, Sparse subspace clustering: Algorithm, theory, and applications, IEEE Transactions on Pattern Analysis and Machine Intelligence.

[6] J. Shi, J. Malik, Normalized cuts and image segmentation, IEEE Transactions on Pattern Analysis and Machine Intelligence 22 (8) (2000) 888–905.

[7] C. Tomasi, T. Kanade, Shape and motion from image streams under orthography: a factorization method, International Journal of Computer Vision 9 (2) (1992) 137–154.

[8] J. P. Costeira, T. Kanade, A multibody factorization method for independently moving objects, International Journal of Computer Vision 29 (3) (1998) 159–179.

[9] K. Kanatant, Motion segmentation by subspace separation: Model selection and reliability evaluation, International Journal of Image and Graphics 2 (2) (2002) 179–197.

[10] B. Jacquet, R. Angst, M. Pollefeys, Articulated and restricted motion subspaces and their signatures, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013, pp. 1506–1513.

[11] R. Basri, D. W. Jacobs, Lambertian reflectance and linear subspaces, IEEE Transactions on Pattern Analysis and Machine Intelligence 25 (2) (2003) 218–233.

[12] A. S. Georghiades, P. N. Belhumeur, D. J. Kriegman, From few to many: Illumination cone models for face recognition under variable lighting and pose, Pattern Analysis and Machine Intelligence, IEEE Transactions on 23 (6) (2001) 643–660.

[13] W. Hong, J. Wright, K. Huang, Y. Ma, Multiscale hybrid linear models for lossy image representation, IEEE Transactions on Image Processing 15 (12) (2006) 3655–3671.

[14] Y. Zhang, Z. Jiang, L. S. Davis, Learning structured low-rank representations for image classification, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013, pp. 676–683.

[15] G. Bull, J. Gao, Transposed low rank representation for image classification, in: International Conference on Digital Image Computing Techniques and Applications, DICTA, Fremantle, Australia, 2012, pp. 1–7. doi:10.1109/DICTA.2012.6411718.

[16] R. Liu, Z. Lin, F. De la Torre, Z. Su, Fixed-rank representation for unsupervised visual learning, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2012, pp. 598–605.

[17] G. Liu, S. Yan, Latent low-rank representation for subspace segmentation and feature extraction, in: IEEE International Conference on Computer Vision (ICCV), 2011, pp. 1615–1622.

[18] A. Y. Yang, J. Wright, Y. Ma, S. S. Sastry, Unsupervised segmentation of natural images via lossy data compression, Computer Vision and Image Understanding 110 (2) (2008) 212–225.

[19] B. Cheng, G. Liu, J. Wang, Z. Huang, S. Yan, Multi-task low-rank affinity pursuit for image segmentation, in: IEEE International Conference on Computer Vision (ICCV), 2011, pp. 2439–2446.

[20] Y. Zhu, D. Huang, F. D. L. Torre, S. Lucey, Complex non-rigid motion 3d reconstruction by union of subspaces, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014, pp. 1542–1549.

[21] S. Tierney, J. Gao, Y. Guo, Subspace clustering for sequential data, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014, pp. 1019–1026.

[22] R. Vidal, Y. Ma, S. Sastry, Generalized principal component analysis (GPCA), IEEE Transactions on Pattern Analysis and Machine Intelligence 27 (12) (2005) 1945–1959.

[23] Y. Guo, J. Gao, F. Li, S. Tierney, M. Yin, Low rank sequential subspace clustering, in: Neural Networks (IJCNN), 2015 International Joint Conference on, IEEE, 2015, pp. 1–8.

[24] Y. Ma, A. Y. Yang, H. Derksen, R. Fossum, Estimation of subspace arrangements with applications in modeling and segmenting mixed data, SIAM review 50 (3) (2008) 413–458.

[25] M. E. Tipping, C. M. Bishop, Mixtures of probabilistic principal component analyzers, Neural computation 11 (2) (1999) 443–482.

[26] Y. Sugaya, K. Kanatani, Geometric structure of degeneracy for multi-body motion segmentation, in: Statistical Methods in Video Processing, Springer, 2004, pp. 13–25.

[27] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, Y. Ma, Robust recovery of subspace structures by low-rank representation, IEEE Transactions on Pattern Analysis and Machine Intelligence 35 (1) (2013) 171–184. doi:10.1109/TPAMI.2012.88.

[28] X. Peng, L. Zhang, Y. Zhang, Constructing l2-graph for subspace learning and segmentation, arXiv:1209.0841v5.

[29] R. Vidal, P. Favaro, Low rank subspace clustering (LRSC), Pattern Recognition Letters 43 (0) (2014) 47 – 61, {ICPR2012} Awarded Papers. doi:http://dx.doi.org/10.1016/j.patrec.2013.08.006.

[30] Y. Guo, J. Gao, F. Li, Spatial subspace clustering for drill hole spectral data, Journal of Applied Remote Sensing 8 (1) (2014) 1–19.

[31] M. Soltanolkotabi, E. Elhamifar, E. J. Candes, et al., Robust subspace clustering, The Annals of Statistics 42 (2) (2014) 669–699.

[32] X. Peng, L. Zhang, Z. Yi, Scalable sparse subspace clustering, in: Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on, IEEE, 2013, pp. 430–437.

[33] J. A. Tropp, A. C. Gilbert, Signal recovery from random measurements via orthogonal matching pursuit, IEEE Transactions on information theory 53 (12) (2007) 4655–4666.

[34] E. L. Dyer, A. C. Sankaranarayanan, R. G. Baraniuk, Greedy feature selection for subspace clustering, The Journal of Machine Learning Research 14 (1) (2013) 2487–2517.

[35] C. You, D. Robinson, R. Vidal, Scalable sparse subspace clustering by orthogonal matching pursuit, in: IEEE Conference on Computer Vision and Pattern Recognition, Vol. 1, 2016.

[36] K. B. Petersen, M. S. Pedersen, et al., The matrix cookbook, Technical University of Denmark 7 (2008) 15.

[37] B. L. Sturm, M. G. Christensen, Comparison of orthogonal matching pursuit implementations., in: EUSIPCO, 2012, pp. 220–224.

[38] M. Yaghoobi, D. Wu, M. E. Davies, Fast non-negative orthogonal matching pursuit, IEEE Signal Processing Letters 22 (9) (2015) 1229–1233.

[39] B. Mailhé, R. Gribonval, P. Vandergheynst, F. Bimbot, Fast orthogonal sparse approximation algorithms over local dictionaries, Signal Processing 91 (12) (2011) 2822–2835.

[40] D. Park, C. Caramanis, S. Sanghavi, Greedy subspace clustering, in: Advances in Neural Information Processing Systems, 2014, pp. 2753–2761.

[41] C. You, C.-G. Li, D. P. Robinson, R. Vidal, Oracle based active set algorithm for scalable elastic net subspace clustering, in: IEEE Conference on Computer Vision and Pattern Recognition, Vol. 1, 2016.

[42] R. Heckel, H. Bölcskei, Robust subspace clustering via thresholding, arXiv preprint arXiv:1307.4891.

[43] X. Zhang, F. Sun, G. Liu, Y. Ma, Fast low-rank subspace segmentation, IEEE Transactions on Knowledge and Data Engineering 26 (5) (2014) 1293–1297. doi:10.1109/TKDE.2013.114.

[44] A. Beck, M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, SIAM Journal of Imaging Sciences 2 (1) (2009) 183–202.

[45] S. Ji, J. Ye, An accelerated gradient method for trace norm minimization, in: Proceedings of the 26th annual international conference on machine learning, ACM, 2009, pp. 457–464.

[46] F. Bach, R. Jenatton, J. Mairal, G. Obozinski, Convex optimization with sparsity-inducing norms, Optimization for Machine Learning (2011) 19–53.

[47] J. Liu, J. Ye, Efficient l1/lq norm regularization, arXiv preprint arXiv:1009.4766.

[48] C. Fowlkes, S. Belongie, F. Chung, J. Malik, Spectral grouping using the Nystrom method, IEEE Transactions on Pattern Analysis and Machine Intelligence 26 (2) (2004) 214–225.

[49] L. Zelnik-Manor, P. Perona, Self-tuning spectral clustering, in: Advances in neural information processing systems, 2004, pp. 1601–1608.

[50] R. Vidal, A tutorial on subspace clustering, IEEE Signal Processing Magazine 28 (2) (2011) 52–68.

[51] M. Muja, D. G. Lowe, Fast approximate nearest neighbors with automatic algorithm configuration, in: International Conference on Computer Vision Theory and Application VISSAPP'09), INSTICC Press, 2009, pp. 331–340.

[52] K. He, J. Sun, Computing nearest-neighbor fields via propagation-assisted kd-trees, in: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE, 2012, pp. 111–118.

[53] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, A. Y. Wu, An optimal algorithm for approximate nearest neighbor searching fixed dimensions, Journal of the ACM (JACM) 45 (6) (1998) 891–923.

[54] M. Muja, D. G. Lowe, Scalable nearest neighbor algorithms for high dimensional data, IEEE Transactions on Pattern Analysis and Machine Intelligence 36.

[55] M. Soltanolkotabi, E. J. Candes, et al., A geometric analysis of subspace clustering with outliers, The Annals of Statistics 40 (4) (2012) 2195–2238.

[56] R. Tron, R. Vidal, A benchmark for the comparison of 3-d motion segmentation algorithms, in: Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, IEEE, 2007, pp. 1–8.

[57] N. . Berestycki, R. Nickl, Concentration of measure, On their websites (2009).

[58] R. Vershynin, Introduction to the non-asymptotic analysis of random matrices, arXiv preprint arXiv:1011.3027.

[59] R. Liu, Z. Lin, Z. Su, Linearized alternating direction method with parallel splitting and adaptive penalty for separable convex programs in machine learning, in: ACML, 2013, pp. 116–132.

Appendix A. Analysis of True Discoveries by kNN

We need the concentration of measure Theorem 1 in [57]

Theorem Appendix A.1. *Suppose that $Z \in \mathbb{R}^d$ is uniformly distributed on unit sphere $\mathbb{S}^{d-1}$ and A is any subset with $P(A) \leq 1/2$. A_ϵ is the ϵ-neighborhood as*

$$A_\epsilon = \{\mathbf{x} \in \mathbb{S}^{d-1} : d(\mathbf{x}, \mathbf{y}) < \epsilon \text{ for some } \mathbf{y} \in A\}$$

then Z belongs to A_ϵ with some probability, i.e.

$$P(Z \in A_\epsilon) \leq 1 - e^{-(d-2)\epsilon^2/2}.$$

Then we obtain the concentration of samples in a small patch when the number of samples is large.

Theorem Appendix A.2. *Let A_ϵ be the small ϵ-neighborhood defined previously in $\mathbb{R}^d$ on $\mathbb{S}^{d-1}$, and N the number of random variables uniformly distributed on $\mathbb{S}^{d-1}$. For $N = \frac{k_0}{\epsilon^2} e^{(d-2)\epsilon^2/2}$, we have*

$$P(K > k_0/C) \geq 1 - e^{-k_0}(eC)^{k_0/C}.$$

where k_0 and $C > 1$ are some constant.

Proof. For any given $\mathbf{x} \in \mathbb{R}^d$, using Theorem Appendix A.1, we know that

$$p = P(\mathbf{x} \in A_\epsilon) \leq 1 - e^{-(d-2)\epsilon^2/2}.$$

Then the event that K variables fall into A_ϵ follows binomial distribution $\mathrm{B}(N, p)$

$$P(K = k) = \binom{N}{k} p^k (1-p)^{N-k}.$$

When N is large and p is small, which is the case here, the above binomial distribution can be approximated by a Poisson distribution $\mathrm{Poi}(\lambda)$ with

$$\lambda = N(1 - e^{-(d-2)\epsilon^2/2}).$$

Using Chernoff bound on Poisson distribution

$$P(K \leq k) \leq e^{-\lambda+k}(\frac{\lambda}{e})^k.$$

Substitute the condition on N, we have

$$P(K \leq \lambda/C) \leq e^{-k_0}(eC)^{k_0/C}.$$

This completes the proof. □

The above theorem states that if N is large enough, with large probability there are k_0/C variables in side A_ϵ.

Next we bound the inner product between samples from different subspaces. Here we use the arguments in Lemma 7.5 in [55].

Lemma Appendix A.3. *Let $\mathbf{A} \in \mathbb{R}^{d_1 \times N_1}$ be a matrix with columns uniformly distributed in $\mathbb{S}^{d_1-1}$, $\mathbf{y} \in \mathbb{R}^{d_2}$ be a vector uniformly sampled from $\mathbb{S}^{d_2-1}$ and a*

deterministic matrix $\mathbf{\Sigma} \in \mathbb{R}^{d_1 \times d_2}$. *For* $t > 0 \in \mathbb{R}$, *the inner product between any column in* $\mathbf{A}$ *and* $\mathbf{\Sigma y}$ *is bounded as follows*

$$\mathbf{a}_i^\top \mathbf{\Sigma y} \leq \frac{2\sqrt{t \log N_1 + t^2}\|\mathbf{\Sigma}\|_F}{\sqrt{d_1}}$$

with probability at least $1 - 2e^{-t}$.

Proof. Using Borell's inequality on the mapping $\mathbf{y} \mapsto \|\mathbf{\Sigma y}\|$ with Lipschitz constant of σ_1, the largest singular value of $\mathbf{\Sigma}$ leads to

$$P(\|\mathbf{\Sigma y}\| > \varepsilon + \sqrt{E\|\mathbf{\Sigma y}\|^2}) < e^{-\frac{1}{2}\varepsilon^2/\sigma_1^2}.$$

As $E\|\mathbf{\Sigma y}\|^2 = \|\mathbf{\Sigma}\|_F^2/d_2$, we choose $\varepsilon = (b-1)\|\mathbf{\Sigma}\|_F/\sqrt{d_2}$ so that

$$P(\|\mathbf{\Sigma y}\| > \frac{b\|\mathbf{\Sigma}\|_F}{\sqrt{d_2}}) < e^{-\frac{1}{2}(b-1)^2/d_2}, \tag{A.1}$$

where we used the fact that $\|\mathbf{\Sigma}\|_F/\sigma_1 > 1$.

The next step is to bound the inner product of a column in $\mathbf{A}$ ($\mathbf{a}_i$ $i = 1 \ldots N_1$) with any vector $\mathbf{x} \in \mathbb{R}^{d_1}$ by upper bound of spherical caps

$$P(\mathbf{a}^\top \mathbf{x} > \varepsilon\|\mathbf{x}\|) < e^{-\frac{1}{2}d_1\varepsilon^2},$$

which leads to the following using the union bound

$$P(\bigcup_i \mathbf{a}_i^\top \mathbf{x} > \varepsilon\|\mathbf{x}\|) < N_1 e^{-\frac{1}{2}d_1\varepsilon^2}. \tag{A.2}$$

Let $\varepsilon = \sqrt{2\log N_1 + 2t}$, $b = \sqrt{2d_2 t}$. Substituting (A.1) to (A.2) gives

$$P(\bigcup_i \mathbf{a}_i^\top \mathbf{\Sigma y} > \frac{2\sqrt{t\log N_1 + t^2}\|\mathbf{\Sigma}\|_F}{\sqrt{d_1}}) \leq 2e^{-t}.$$

Therefore

$$P(\bigcap_i \mathbf{a}_i^\top \mathbf{\Sigma y} \leq \frac{2\sqrt{t\log N_1 + t^2}\|\mathbf{\Sigma}\|_F}{\sqrt{d_1}}) \geq 1 - 2e^{-t},$$

which concludes the proof. □

The above gives the upper bound of the inner product, which connects to samples in subspaces with the following corollary.

Corollary Appendix A.4. *Let $\mathbf{X} \in \mathbb{R}^{d\times N_1}$ be a matrix with columns randomly sampled from subspace $\mathcal{S}_1$ with dimensionality d_1 and $\mathbf{y} \in \mathbb{R}^d$ be a sample from subspace $\mathcal{S}_2$ with dimensionality d_2. Both the $\mathbf{y}$ and columns of $\mathbf{X}$ have been rescaled to have unitary ℓ_2 norm. The inner product between any column in $\mathbf{X}$ and $\mathbf{y}$ is bounded as the following*

$$\mathbf{x}_i^\top \mathbf{y} \leq \frac{2\mathcal{A}_{1,2}\sqrt{\min\{d_1, d_2\}(t\log N_1 + t^2)}}{\sqrt{d_1}}$$

with probability at least $1 - 2e^{-t}$ for t given previously.

This is the simple application of Appendix A.3 with $\mathbf{\Sigma} = \mathbf{U}_1^\top \mathbf{U}_2$ where $\mathbf{U}_j$ $j = 1, 2$ is the orthonormal basis for subspace $\mathcal{S}_j$ and $\mathcal{A}_{1,2}$ is the affinity between subspaces $\mathcal{S}_1$ and $\mathcal{S}_2$ described in Definition 1.2 in [31], which we recall here:

$$\mathcal{A}_{\mathcal{S}_i,\mathcal{S}_j} = \sqrt{\frac{\cos^2\theta^{(1)} + \cdots + \cos^2\theta^{(d_i\wedge d_j)}}{d_i \wedge d_j}}$$

where $\{\cos^2\theta^{(1)}, \ldots, \cos^2\theta^{(d_i\wedge d_j)}\}$ are the principal angles between subspaces $\mathcal{S}_i$ and $\mathcal{S}_j$. Please see Definition 1.1 in [31] for full details.

Without loss generality, we consider a sample $\mathbf{x}_1$ from subspace $\mathcal{S}_1$. The following theorem ensures that kNN will find k nearest neighbors of $\mathbf{x}_1$ from $\mathcal{S}_1$ only.

Theorem Appendix A.5. *Let d_m be the minimum dimensionality of all the subspaces. Given the conditions in Theorem Appendix A.2, if $\mathcal{A}_{\ell,1} \leq \frac{\sqrt{d_m}(1-\epsilon^2/2)}{2\sqrt{d(t\log N_\ell + t^2)}}$, then the samples selected for any sample $\mathbf{x}_1$ from subspace $\mathcal{S}_1$ by using kNN (with $k = k_0/C$) contains no samples from other subspaces but $\mathcal{S}_1$ with probability at least $1 - 2e^{-t} - e^{-k_0}(eC)^{k_0/C}$.*

Proof. This is a straightforward application of Theorem Appendix A.2 and Collary Appendix A.4 and the following. If

$$A_{\ell,1} \leq \frac{\sqrt{d_\ell}(1-\epsilon^2/2)}{2\sqrt{\min\{d_1, d_2\}(t\log N_\ell + t^2)}}$$

then

$$\frac{2\mathcal{A}_{1,2}\sqrt{\min\{d_1, d_2\}(t\log N_1 + t^2)}}{\sqrt{d_1}} \leq 1 - \epsilon^2/2.$$

The distance and inner product is connected by

$$d(\mathbf{x}, \mathbf{y}) = 2 - 2\mathbf{x}^\top \mathbf{y},$$

when vectors are normalised. □

The above discussion deals with clean data only. In the following section we show that the results are similar for noisey data as long as the noise level is not too great. We begin with the following series of lemmas

Lemma Appendix A.6. *Let random variables X and Y in $\mathbb{R}^d$ both be from Gaussian distribution $\mathcal{N}(0, \sigma^2\mathbf{I})$. For any given positive δ, we have*

$$P(|X^\top Y| > \delta) \leq \frac{d\sigma^4}{\delta^2}.$$

Proof. First we assume X and Y are standard Gaussian, we have

$$P(|X^\top Y| > \delta) = P((X^\top Y)^2 > \delta^2) \leq \frac{E(X^\top Y)^2}{\delta^2},$$

where the inequality is by Chernoff bound. Since both X and Y are both standard Gaussian, they are isotropic, so we have

$$E(X^\top Y)^2 = d.$$

The above can be obtained by

$$E(X^\top Y)^2 = E_X\{E_{X|Y}(X^\top y)^2\} = E_Y(|y|^2) = d,$$

where the first equality comes from law of iterative expectation, the second from isotropic property and the last from the fact that sum of standard Gaussian is Chi-square with d degrees of freedom.

After proper rescaling, we obtain the result in the lemma. □

Lemma Appendix A.7. *Let $X \in \mathbb{R}^d$ be Gaussian random variable from $\mathcal{N}(0, \sigma^2\mathbf{I})$ and $\mathbf{y} \in \mathbb{R}^d$ be a fixed vector. The following holds with any positive δ*

$$p(|X^\top \mathbf{y}| \geq \delta) \leq \exp(1 - \frac{c\delta^2}{\sigma^2\|\mathbf{y}\|_2^2}),$$

where c is a constant related to sub-Gaussian norm [58] of a standard Gaussian.

This is a straightforward application of sub-Gaussian tail to $X^\top \mathbf{y}$ with rescaling. Now we consider the inner product between two unitary vectors in subspaces with noise. We use the following model

$$\mathbf{Y} = \mathbf{X} + \mathbf{E} \tag{A.3}$$

where Y is the observation, X is the clear signal in some subspace and E is the noise assumed to be from $\mathcal{N}(0, \sigma^2\mathbf{I})$. We assume that the observations have been rescaled properly such that X is from a unit sphere in $\mathbb{S}^{d-1}$ and the variance of the noise is bounded, i.e. $\sigma < \sigma_0$.

First we note that under these conditions, the noise can increase and decrease inner product between observed signals by only a small amount, which is shown in the following lemma.

Lemma Appendix A.8. *Let $\mathbf{y}_i$ $(i = 1, 2)$ be observations from the model in (A.3), such that $\mathbf{y}_i = \mathbf{x}_i + \mathbf{e}_i$, $\|\mathbf{x}_i\|_2 = 1$ and $\mathbf{e}_i \sim \mathcal{N}(0, \sigma^2\mathbf{I})$. If $P(|\mathbf{x}_1^\top \mathbf{x}_2| > v) \geq p$, we have*

$$P(\mathbf{y}_1^\top \mathbf{y}_2 > v - 3\delta) \geq p - \exp(1 - \frac{c\delta^2}{\sigma^2}) - \frac{d\sigma^4}{2\delta^2}.$$

If $P(\mathbf{x}_1^\top \mathbf{x}_2 < v) \geq p$, we have

$$P(\mathbf{y}_1^\top \mathbf{y}_2 < v + 3\delta) \geq p - \exp(1 - \frac{c\delta^2}{\sigma^2}) - \frac{d\sigma^4}{2\delta^2}.$$

Proof. We prove the $P(\mathbf{x}_1^\top \mathbf{x}_2 > v) \geq p$ case. The other cases can be proved similarly. Writing $\mathbf{y}_1^\top \mathbf{y}_2$ in terms and using triangular inequality gives

$$|\mathbf{y}_1^\top \mathbf{y}_2| \geq |\mathbf{x}_1^\top \mathbf{x}_2| - |\mathbf{x}_1^\top \mathbf{e}_2| - |\mathbf{e}_1^\top \mathbf{x}_2| - |\mathbf{e}_1^\top \mathbf{e}_2|.$$

Then

$$P(\mathbf{y}_1^\top \mathbf{y}_2 > v - 3\delta) = P(|\mathbf{x}_1^\top \mathbf{x}_2| > v \bigcap \mathbf{x}_1^\top \mathbf{e}_2 \geq -\delta \bigcap \mathbf{e}_1^\top \mathbf{x}_2 \geq -\delta \bigcap \mathbf{e}_1^\top \mathbf{e}_2 \geq -\delta).$$

By using Lemma Appendix A.6 and Appendix A.7, we obtain the desired result. □

Lemma Appendix A.8 states that the noise will dispel the vectors when they are very close and attract them when they are far away in terms of the inner product induced distance. The effect of noise for a given sample in subspace $\mathcal{S}_1$ is then to make the samples from other subspaces closer to it and more difficult to separate reflected by the reduced probability as shown in the following theorem.

Theorem Appendix A.9. *Let d_m be the minimum dimensionality of all the subspaces. Given the conditions in Theorem Appendix A.2 and the noise model* (A.3)*, for a small positive δ, if $\mathcal{A}_{\ell,1} \leq \frac{\sqrt{d_m}(1-\epsilon^2/2-6\delta)}{2\sqrt{d(t\log N_\ell + t^2)}}$, then the samples selected for any sample $\mathbf{y}_1$ from subspace $\mathcal{S}_1$ by using kNN (with $k = k_0/C$) contains no samples from other subspaces but $\mathcal{S}_1$ with probability at least $1 - 2e^{-t} - e^{-k_0}(eC)^{k_0/C} - 2\exp(1 - \frac{c\delta^2}{\sigma^2}) - \frac{d\sigma^4}{\delta^2}$.*

Proof. According to the model (A.3), $\mathbf{y}_1 = \mathbf{x}_1 + \mathbf{e}_1$ and $\mathbf{x}_1$ is in a unit sphere. From Theorem Appendix A.2, we know that there are at least $\lfloor k_0/C \rfloor$ samples from $\mathcal{S}_1$ in the patch A_ϵ centred at $\mathbf{x}_1$ with probability at least $1-e^{-k_0}(eC)^{k_0/C}$. This leads to the following combing Lemma Appendix A.8

$$P(\min_{j\in\mathcal{N}_1}\{\mathbf{y}_1^\top \mathbf{y}_{i_j}\} \geq 1 - \epsilon^2/2 - 3\delta) \geq 1 - e^{-k_0}(eC)^{k_0/C} - \exp(1 - \frac{c\delta^2}{\sigma^2}) - \frac{d\sigma^4}{2\delta^2}$$

where $\mathcal{N}_1$ is the set of $\lfloor k_0/C \rfloor$ samples around $\mathbf{x}_1$ in A_ϵ patch.

Using Corollary Appendix A.4 and Lemma Appendix A.8 results in that with probability at least $1 - 2e^{-t} - \exp(1 - \frac{c\delta^2}{\sigma^2}) - \frac{d\sigma^4}{2\delta^2}$

$$\mathbf{y}_i^\top \mathbf{y}_1 \leq \frac{2\mathcal{A}_{1,2}\sqrt{\min\{d_1, d_2\}(t\log N_1 + t^2)}}{\sqrt{d_1}} + 3\delta$$

for any $\mathbf{y}_i$ from subspace $\mathcal{S}_\ell$.

Similar to Theorem Appendix A.5, combing the above two statements, if

$$A_{\ell,1} \leq \frac{\sqrt{d_\ell}(1 - \epsilon^2/2 - 6\delta)}{2\sqrt{\min\{d_1, d_2\}(t\log N_\ell + t^2)}}$$

then

$$\frac{2\mathcal{A}_{1,2}\sqrt{\min\{d_1, d_2\}(t\log N_1 + t^2)}}{\sqrt{d_1}} + 3\delta \leq 1 - \epsilon^2/2 - 3\delta$$

with the probability stated in this theorem. □

To obtain the high probability, it is required that σ be as low as possible given that δ is at the scale of ϵ and the probability is also tied up with the ambient dimensionality.

Appendix B. Solving Exact kSSC via LADMPSAP

The exact objective for kSSC is as follows:

$$\min_{\mathbf{Z},\mathbf{E}} \lambda\|\mathbf{Z}\|_1 + \frac{1}{2}\|\mathbf{E}\|_F^2 \quad \text{s.t. } \mathbf{x}_i = \sum_{j\in\Omega_i}^{k} \mathbf{x}_j Z_{ji} + \mathbf{e}_i, \operatorname{diag}(\mathbf{Z}) = \mathbf{0}. \tag{B.1}$$

To solve (B.1) one can use use LADMPSAP (Linearized Alternating Direction Method with Parallel Spliting and Adaptive Penalty) [59]. Essentially LADMPSAP converts the original objective into a series of closed form proximity problems which can be solved at an element wise level. This allows us to resolve the first constraint of (B.1) since we can enforce it by ignoring those elements outside of Ω. LADMPSAP provides guaranteed convergence and we refer readers to [59] for full details. Furthermore the sub-steps involving the update of primary variables are independent of each other and if one chooses can be computer parallel. However since we already compute each column of $\mathbf{Z}$ in parallel we find that such further parallelisation is unnecessary.

The full algorithm can be found in Algorithm 3. We begin discussing the details by re-writing, with some abuse of notation, the original objective (B.1) for a single column of $\mathbf{Z}$ and $\mathbf{E}$

$$\min_{\mathbf{z}_i,\mathbf{e}_i} \lambda\|\mathbf{z}_i\|_1 + \frac{1}{2}\|\mathbf{e}_i\|_2^2 \quad \text{s.t. } \mathbf{x}_i = \mathbf{X}_i\mathbf{z}_i + \mathbf{e}_i \tag{B.2}$$

where $\mathbf{z}_i = \mathbf{z}_{\Omega_i i}$ i.e. the vector of rows Ω_i and column i of $\mathbf{Z}$, $\mathbf{e}_i$ is column i of $\mathbf{E}$ and $\mathbf{X}_i = \mathbf{X}_{(:,\Omega_i)}$ i.e. the matrix formed from the columns of $\mathbf{X}$ indexed by Ω_i. Note that we have removed the constraint $\operatorname{diag}(\mathbf{Z}) = \mathbf{0}$ since we enforce it by ensuring that no diagonal entries are present in each Ω_i.

Next we relax the constraints and form the Lagrangian function

$$\min_{\mathbf{z}_i,\mathbf{e}_i,\mathbf{y}_i} \lambda\|\mathbf{z}_i\|_1 + \frac{1}{2}\|\mathbf{e}_i\|_F^2 + \langle \mathbf{y}_i, \mathbf{X}_i\mathbf{z}_i - \mathbf{x}_i + \mathbf{e}_i\rangle + \frac{1}{2}\|\mathbf{X}_i\mathbf{z}_i - \mathbf{x}_i + \mathbf{e}_i\|_2^2 \tag{B.3}$$

where $\mathbf{y}_i$ is a Lagrange multiplier. Then LADMPSAP consists of an iterative procedure over the optimisation variables We outline the procedure in Algorithm 3 Here we detail the solutions to the sub problems. Denote $\mathbf{z}_i^{\mathtt{t}}, \mathbf{e}_i^{\mathtt{t}}, \mathbf{y}_i^{\mathtt{t}}$ the variables at iteration t. Then:

1. Fix others and update $\mathbf{z}_i^{\mathtt{t}+1}$

$$\mathbf{z}_i^{\mathtt{t}+1} = \underset{\mathbf{z_i}}{\operatorname{argmin}} \lambda\|\mathbf{z}_i\|_1 + \frac{\rho}{2}\|\mathbf{z}_i - (\mathbf{z}_i^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{z}_i^{\mathtt{t}}))\|_2^2$$

   where $\partial F(\mathbf{z}_i^{\mathtt{t}}) = \mu^{\mathtt{t}}\mathbf{X}_i^T(\mathbf{X}_i\mathbf{z}_i^{\mathtt{t}} - (\mathbf{x}_i - \mathbf{e}_i^{\mathtt{t}} - \frac{1}{\mu}\mathbf{y}_i^{\mathtt{t}}))$. For which we have a closed form solution given by

$$\mathbf{z}_i^{\mathtt{t}+1} = \operatorname{sign}(\mathbf{b}^{\mathtt{t}})\max(|\mathbf{b}^{\mathtt{t}}| - \frac{\lambda}{\rho}, 0), \tag{B.4}$$

   where $\mathbf{b}_i^{\mathtt{t}} = \mathbf{z}_i^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{z}_i^{\mathtt{t}})$, see [46, 47] for details.

2. Fix others and update $\mathbf{e}_i^{\mathtt{t}+1}$

$$\mathbf{e}^{\mathtt{t}+1} = \underset{\mathbf{e_i}}{\operatorname{argmin}} \frac{1}{2}\|\mathbf{e}_i\|_2^2 + \frac{\mu^t}{2}\|\mathbf{e}_i - (\mathbf{x}_i - \mathbf{X}_i\mathbf{z}_i^{\mathtt{t}} - \frac{1}{\mu^t}\mathbf{y}_i^{\mathtt{t}})\|_2^2$$

   which has the closed form solution given by

$$\mathbf{e}_i^{\mathtt{t}+1} = \frac{\mathbf{x}_i - \mathbf{X}_i\mathbf{z}_i^{\mathtt{t}} - \frac{1}{\mu^{\mathtt{t}}}\mathbf{y}_i^{\mathtt{t}}}{\frac{1}{\mu^{\mathtt{t}}} + 1}, \tag{B.5}$$

3. Update $\mathbf{y}_i^{\mathtt{t}+1}$

$$\mathbf{y}_i^{\mathtt{t}+1} = \mathbf{y}_i^{\mathtt{t}} + \mu^{\mathtt{t}}(\mathbf{X}_i\mathbf{z}_i^{\mathtt{t}+1} - \mathbf{x}_i + \mathbf{e}_i^{\mathtt{t}+1}) \tag{B.6}$$

Appendix C. Solving Relaxed SSC via Accelerated Gradient Descent

Here we discuss how to solve (4) via FISTA. The full algorithm is outlined in Algorithm 4. Denote

$$\min_{\mathbf{Z}} L = \lambda\|\mathbf{Z}\|_1 + \frac{1}{2}\|\mathbf{X} - \mathbf{X}\mathbf{Z}\|_F^2 \tag{C.1}$$

and

$$\min_{\mathbf{Z}} \widetilde{L}_\rho(\mathbf{Z}, \mathbf{Z}^{\mathtt{t}}) = \lambda\|\mathbf{Z}\|_1 + \frac{\rho}{2}\|\mathbf{Z} - (\mathbf{Z}^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{Z}^{\mathtt{t}}))\|_F^2,$$

Algorithm 3 Solving (B.1) via LADMPSAP

Require: $\mathbf{x}_i$, $\mathbf{X}_i$, Ω_i, λ, ϵ_1, ϵ_2,

1: Initialise: $\mathbf{z}_i = \mathbf{0}$, $\mathbf{e}_i = \mathbf{0}$, $\mathbf{y}_i = \mathbf{0}$, $\mu = 0.1$, $\mu_{\max} = 1$, $\gamma^0 = 1.1$, $\rho = \|\mathbf{X}_i\|_F$

2: **while** not converged **do**

3: Update $\mathbf{z}_i^{\mathtt{t}+1}$ using (B.4)

4: Update $\mathbf{e}_i^{\mathtt{t}+1}$ using (B.5)

5: Set q

$$q = \frac{\mu^{\mathtt{t}}\sqrt{\rho}}{\|\mathbf{X}_i\|_F}\max(\|\mathbf{z}_i^{\mathtt{t}+1} - \mathbf{z}_i^{\mathtt{t}}\|_2, \|\mathbf{e}_i^{\mathtt{t}+1} - \mathbf{e}_i^{\mathtt{t}}\|_2)$$

6: Check stopping criteria

$$\frac{\|\mathbf{X}_i\mathbf{z}_i^{\mathtt{t}+1} - \mathbf{x}_i + \mathbf{e}_i^{\mathtt{t}+1}\|_2}{\|\mathbf{X}_i\|_F} < \epsilon_1, q < \epsilon_2$$

7: Update $\mathbf{y}_i^{\mathtt{t}+1}$ using (B.6)

8: Update $\gamma^{\mathtt{t}+1}$

$$\gamma^{\mathtt{t}+1} = \begin{cases} \gamma^0 & \text{if } q < \epsilon_2 \\ 1 & \text{otherwise,} \end{cases}$$

9: Update $\mu^{\mathtt{t}+1}$

$$\mu^{\mathtt{t}+1} = \min(\mu_{\max}, \gamma\mu^{\mathtt{t}})$$

10: **end while**

11: **return** $\mathbf{z}_i$

$F = \frac{1}{2}\|\mathbf{X} - \mathbf{X}\mathbf{Z}\|_F^2$ and $\partial F = -\mathbf{X}^T(\mathbf{X} - \mathbf{X}\mathbf{Z}^{\mathtt{t}})$. The solution is given by the closed-form ℓ_1 shrinkage function $\mathcal{S}_\tau$ as follows

$$\mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{Z}^{\mathtt{t}}) = \text{sign}(\mathbf{Z}^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{Z}^{\mathtt{t}}))\max(|\mathbf{Z}^t - \frac{1}{\rho}\partial F(\mathbf{Z}^{\mathtt{t}})| - \frac{\lambda}{\rho}, 0). \tag{C.2}$$

Algorithm 4 Solving (4) via FISTA

Require: $r = \infty$, $\mathbf{Z} = \mathbf{0}$, $\mathbf{J} = \mathbf{0}$, $\alpha = 1$, λ, ρ, γ, ϵ

while $r^{\mathtt{t}} - r^{\mathtt{t}-1} \geq \epsilon$ **do**

while $L(\mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{J}^{\mathtt{t}})) \geq \widetilde{L}_{\rho}(\mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{J}^{\mathtt{t}}), \mathbf{J}^{\mathtt{t}})$ **do**

$\rho = \gamma\rho$

end while

$\mathbf{Z}^{\mathtt{t}+1} = \mathcal{S}_{\frac{\lambda}{\rho}}(\mathbf{J}^{\mathtt{t}}))$

$\alpha^{\mathtt{t}+1} = \frac{1+\sqrt{1+4\alpha^{t2})}}{2}$

$\mathbf{J}^{\mathtt{t}+1} = \mathbf{Z}^{\mathtt{t}+1} + \left(\frac{\alpha^{\mathtt{t}}-1}{\alpha^{\mathtt{t}+1}}\right)(\mathbf{Z}^{\mathtt{t}+1} - \mathbf{Z}^{\mathtt{t}})$

$r^{\mathtt{t}+1} = \frac{1}{2}\|\mathbf{X} - \mathbf{X}\mathbf{Z}^{\mathtt{t}+1}\|_F^2 + \lambda\|\mathbf{Z}^{\mathtt{t}+1}\|_1$

end while

Appendix D. Solving Exact SSC via LADMPSAP

Here we discuss how to solve (5) via LADMPSAP. The full algorithm is outlined in Algorithm 5. Denote the Augmented Lagrangian form

$$\min_{\mathbf{E},\mathbf{Z}} \frac{1}{2}\|\mathbf{E}\|_F^2 + \lambda\|\mathbf{Z}\|_1 + \langle \mathbf{Y}, \mathbf{X}\mathbf{Z} - \mathbf{X} + \mathbf{E}\rangle + \frac{\mu}{2}\|\mathbf{X}\mathbf{Z} - \mathbf{X} + \mathbf{E}\|_F^2$$

Iterate the following:

1. Fix others and update $\mathbf{Z}^{\mathtt{t}+1}$

$$\min_{\mathbf{Z}} \lambda\|\mathbf{Z}\|_1 + \langle \mathbf{Y}, \mathbf{X}\mathbf{Z}\rangle + \frac{\mu}{2}\|\mathbf{X}\mathbf{Z} - (\mathbf{X} - \mathbf{E})\|_F^2$$

$$\min_{\mathbf{Z}} \lambda\|\mathbf{Z}\|_1 + \frac{\rho}{2}\|\mathbf{Z} - (\mathbf{Z}^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{Z}^{\mathtt{t}})\|_F^2$$

where $F = \frac{\mu}{2}\|\mathbf{X}\mathbf{Z} - (\mathbf{X} - \mathbf{E} - \frac{1}{\mu}\mathbf{Y})\|_F^2$ and $\partial F = \mu^{\mathtt{t}}\mathbf{X}^T(\mathbf{X}\mathbf{Z}^{\mathtt{t}} - (\mathbf{X} - \mathbf{E} - \frac{1}{\mu^{\mathtt{t}}}\mathbf{Y}\mathbf{Z}^{\mathtt{t}}))$. Then

$$\mathbf{Z}^{\mathtt{t}+1} = \mathrm{sign}(\mathbf{B}^{\mathtt{t}})\max(|\mathbf{B}^{\mathtt{t}}| - \frac{\lambda}{\rho}, 0), \tag{D.1}$$

where $\mathbf{B}_i^{\mathtt{t}} = \mathbf{Z}^{\mathtt{t}} - \frac{1}{\rho}\partial F(\mathbf{Z}^{\mathtt{t}})$.

2. Fix others and update $\mathbf{E}^{\mathtt{t}+1}$

$$\min_{\mathbf{E}} \frac{1}{2}\|\mathbf{E}\|_F^2 + \langle \mathbf{Y}, \mathbf{E}\rangle + \frac{\mu^{\mathtt{t}}}{2}\|\mathbf{X}\mathbf{Z} - \mathbf{X} + \mathbf{E}\|_F^2$$

$$\min_{\mathbf{E}} \frac{1}{2}\|\mathbf{E}\|_F^2 + \frac{\mu^{\mathfrak{t}}}{2}\|\mathbf{E} - (\mathbf{X} - \mathbf{X}\mathbf{Z}^{\mathfrak{t}} - \frac{1}{\mu^{\mathfrak{t}}}\mathbf{Y}^{\mathfrak{t}})\|_F^2$$

$$\mathbf{e}_i^{\mathfrak{t}+1} = \frac{\mathbf{X} - \mathbf{X}\mathbf{Z}^{\mathfrak{t}} - \frac{1}{\mu^{\mathfrak{t}}}\mathbf{Y}^{\mathfrak{t}}}{\frac{1}{\mu^{\mathfrak{t}}} + 1}, \tag{D.2}$$

3. Update $\mathbf{Y}^{\mathfrak{t}+1}$

$$\mathbf{Y} = \mathbf{Y} + \mu^{\mathfrak{t}}(\mathbf{X}\mathbf{Z}^{\mathfrak{t}+1} - \mathbf{X} + \mathbf{E}^{\mathfrak{t}+1}) \tag{D.3}$$

Algorithm 5 Solving (5) via LADMPSAP

Require: $\mathbf{X}$, λ, ϵ_1, ϵ_2,

Initialise: $\mathbf{Z} = \mathbf{0}$, $\mathbf{E} = \mathbf{0}$, $\mathbf{Y} = \mathbf{0}$, $\mu = 0.1$, $\mu_{\max} = 1$, $\gamma^0 = 1.1$, $\rho = \|\mathbf{X}_i\|_F$

while not converged **do**

Update $\mathbf{Z}_i^{\mathtt{t}+1}$ using (D.1)

Update $\mathbf{E}_i^{\mathtt{t}+1}$ using (D.2)

Set q

$$q = \frac{\mu^{\mathtt{t}}\sqrt{\rho}}{\|\mathbf{X}\|_F}\max(\|\mathbf{Z}^{\mathtt{t}+1} - \mathbf{Z}^{\mathtt{t}}\|_F, \|\mathbf{E}^{\mathtt{t}+1} - \mathbf{E}^{\mathtt{t}}\|_F)$$

Check stopping criteria

$$\frac{\|\mathbf{X}\mathbf{Z}^{\mathtt{t}+1} - \mathbf{X} + \mathbf{E}^{\mathtt{t}+1}\|_2}{\|\mathbf{X}\|_F} < \epsilon_1, q < \epsilon_2$$

Update $\mathbf{Y}_i^{\mathtt{t}+1}$ using (D.3)

Update $\gamma^{\mathtt{t}+1}$

$$\gamma^{\mathtt{t}+1} = \begin{cases} \gamma^0 & \text{if } \ q < \epsilon_2 \\ 1 & \text{otherwise,} \end{cases}$$

Update $\mu^{\mathtt{t}+1}$

$$\mu^{\mathtt{t}+1} = \min(\mu_{\max}, \gamma\mu^{\mathtt{t}})$$

end while

return $\mathbf{Z}$